arXiv:1403.7827v1 [physics.soc-ph] 30 Mar 2014

Patterns of success in co-authorship networks are highly sensitive to author disambiguation

David F. Klosik[a], Stefan Bornholdt[a], Marc-Thorsten Hütt[b]

[a] *Institute for Theoretical Physics, University of Bremen, Hochschulring 18, 28359 Bremen, Germany*

[b] *School of Engineering and Science, Jacobs University Bremen, Campusring 1, 28759 Bremen, Germany*

Abstract

Following the work of Krumov et. al (2011) [1] we revisit the question whether the usage of large citation datasets allows for the quantitative assessment of social (by means of co-authorship of publications) influence on the progression of science. Applying a more comprehensive and well-curated dataset containing the publications in the journals of the American Physical Society during the whole 20th century we find that the measure chosen in the original study, a score computed on small induced subgraphs, has to be used with caution, since the obtained results are highly sensitive against the exact implementation of the author disambiguation task.

Keywords:

1. Introduction

Ever since the seminal work of Thomas S. Kuhn [2] it is widely accepted that the institutional process of knowledge production, i.e. science, cannot be fully described in purely logical, content-related terms, but has a significant social aspect to it. However, although the scientific community provides a comprehensive book-keeping of its efforts by citing earlier work in new publications, only recently has this information been made widely accessible in form of electronic data sets. With the aggregated citation information within a set of scientific publications at hand, one might now be able to quantitatively assess the extent to which the social embedding of science influences its structure and progression.

Traditionally, the focus of citation data analysis has been on the single publication level; indeed, the most prominent property of *paper citation networks* in which the vertices represent papers while directed edges represent the citations between the publications has been described by de Solla Price as early

Email addresses: klosik@itp.uni-bremen.de (David F. Klosik), bornholdt@itp.uni-bremen.de (Stefan Bornholdt), m.huett@jacobs-university.de (Marc-Thorsten Hütt)

as 1965 [3]: the number of citations a paper receives (i.e. the corresponding node's in-degree) is broadly distributed, rendering highly cited publications significantly more frequent than they would be if scientists cited earlier work randomly. Similar broad degree distributions have been found in networks describing e.g. technical, social or biological interactions [4] and the question which might be the mutual underlying process led to a now well-studied model class for network growth governed by a *rich-get-richer* effect, commonly referred to as preferential attachment[5, 6], [7–10], [11, 12]. In more recent times, citation data itself has mainly been used to quantify the assessment of scientific research thereby interpreting citations as indicators of impact or assignment of credit. To this end, numerous quantitative measures have been described which range from counting direct citations to considering also indirect citation paths [13, 14] [15, 16], [17] and some of which aim at the scholar-[18, 19] or the journal-level [20]. Additionally, there have also been structural investigations, e.g. regarding the community level [21] or other topological properties like the richness of feed-forward-loops in citation networks [22].

Another line of research has considered the collaboration network that can be constructed from citation data given the authorship metadata. More precisely, in a *co-authorship network* vertices represent authors and are connected by an undirected link if the two corresponding authors have co-authored one ore more papers together. These networks have been investigated about a decade ago [23, 24] with the main findings being the rather broad degree distribution and the strong small-world effect, i.e. the short paths between scholars in the network. There are also few approaches combining both citation and collaboration data, e.g. [25].

In this study we will construct a co-authorship network from a citation dataset (see Materials and Methods section) while the actual citation information is used to assess the success of the resulting links in the network.

While there have been studies about large-scale properties of co-authorship networks like paths lengths or community sizes and, of course, on the single-vertex scale, following Krumov et. al [1] here the focus will be on the intermediate level of small induced subgraphs known from the investigation of network motifs which originated in the biological context [26, 27] and whose application has drawn some considerable criticism since the null-model graph ensemble has to be chosen with great care [28] [29]. Note, however, that throughout this study *score* does not refer to the otherwise commonly used *z-score*, since we are not interested in subgraph frequencies, but to the q_m-value that is defined below. We will focus on the three- and four-node undirected induced subgraphs shown in Fig. 1.

We construct the collaboration network from a citation dataset provided by the American Physical Society (*see Materials and Methods*) by, first, identifying the articles' authors given in the article metadata. All author instances (i.e. names in author lists) are then grouped so that, ideally, all instances corresponding to the same actual scholar are represented by a single vertex in the network. There are different possible implementations of this *author disambiguation task* which has become a object of investigation in its on right [30, 31]

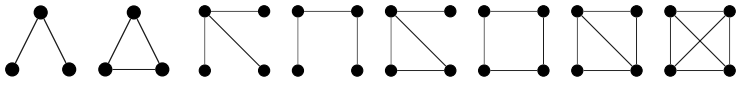

Figure 1: The subgraphs used in this study. In the bottom line the minimum number of distinct papers required to build the subgraph is given.

with methods ranging from directly comparing the author names to incorporating metadata or even citation data. While in [1] one specific version has been used, here we find that the score proposed to assess the correlation between the collaboration pattern and the success of the corresponding papers is very sensitive to the exact implementation of the author disambiguation.

Furthermore, we show how the score is affected by the exclusion of large collaborations when the length of the author list above which papers are discarded is varied; in [1] this value was fixed at a number of eight.

2. Theory/Calculation

First, we review the computation of the score proposed by Krumov et al. Two distinct authors who have co-authored are connected by an edge, e, that represents the list of all their mutually published papers, $P(e)$. Note that by this procedure a single publication can be represented by many edges. In order to assess whether there is a correlation between the collaboration structure on the small subgraph scale and the scientific impact of the publications contained in the subgraph's edges we use the number of citations to those papers to, first, compute the average number of citations to the papers of a single edge

$$w_2 = \langle c \rangle_e = \frac{1}{|P(e)|} \sum_{p \in P(e)} c(p) \tag{1}$$

and then take the average over all edges of all instances of a specific motif, i.e.

$$q_m = \frac{1}{N_m} \sum_{m' \in M_m} q_{m,m'} = \frac{1}{N_m E_m} \sum_{m' \in M_m} \sum_{e \in m'} \langle c \rangle_e. \tag{2}$$

where N_m gives the number of instances of the subgraph of type m and E_m yields the number of its edges (e.g. for the box motif $E_5 = 4$). A sketch of the procedure is given in Fig. 2. If one now shuffles the citation numbers among all publications the q_m is uniformly distributed, indicating that the motif-scale collaboration patterns in the given co-authorship network and the papers' success are not correlated after reshuffling. With the particular choice of (1) the motif scores are expected to yield the average number of citations of all papers in the network for every motif. The fact that this holds in all our computations implies that we obtain proper averages of the often broadly distributed citation numbers. In [1] alternative edge scores have been proposed but with the current

dataset (1) turns out to be the appropriate choice. It is important to notice that the shuffling procedure does only affect the citation data on the edges of the otherwise fixed collaboration network, i.e. no topological shuffling is applied.

Krumov et al. report considerably higher q_m scores for the four-node motif called Motif 5 in Fig. 1 which they call the *box motif*. This subgraph stands out since it needs four distinct publications to be constructed (while e.g. the four-node clique, motif 7, might contain only one single mutually published paper) and there must not be any other collaborations between the four authors than in the author pairs corresponding to the four edges. The box motif therefore is an anti-clustered structure and considerable fewer instances are found than of the other motif types. Again, we stress that we do not focus on motif frequency.

In order to keep the network topology and the dynamical quantity separated, in contrast to [1] we do not discard edges which are composed only of papers that did not receive any citation, but keep them in the network.

2.1. Maximum Collaboration Size

As mentioned above, in [1] all papers with more than eight authors have been discarded. Here we allow for different values of the *maximum collaboration size*, MCS (chapt. 4.1). In terms of the social aspects encoded in a co-authorship network the exclusion of very large collaborations can argued for since the thousands of links between their authors will hardly represent the same degree of personal connection as an edge between e.g. the single two authors of a publication.

2.2. Author Name Disambiguation

While one of the more elaborated author disambiguation schemes might have been used here, in order to assess the sensitivity of the score proposed in [1] we instead chose two versions of a very simple measure which consider only the authors' names as provided in the author lists of the publications in the dateset [32]. In the *all initials* (allINIT) disambiguation two author instance names are considered identical if in addition to the surname all initials are compatible (meaning that J. Smith and John Smith are merged). The second implementation requires the full first name strings to match, we consider this *strict* (STRICT) since it will separate author instances if the first name is abbreviated in one and written out in the other case like the ones from the above example. However, there are surnames common enough to be shared be people with different given names; if these given names have compatible initials the allINIT method will spuriously merge them. In order to address this we apply a third disambiguation which we will refer to as SPLIT: We apply indenpendently the allINIT and STRENG disambiguation to the same data and then track those allINIT-names that are split in more than two distinct authors in the STRENG implementation. All papers with those names among the authors are then filtered. On the resulting dataset we then perform all remaining steps of the calculation with the allINIT method applied. Note that the result of this procedure depends on all previous data filtering, especially

the one according to the maximum collaboration size (MCS), so the number of remaining papers has to be compared to the number of publications in the original data after discarding those with more than MCS authors. In the cases shown here the the data had been prepared with a MCS value of 10 before the application of the SPLIT filtering and we find that while for the merge of the journals 36% and in PRC 32% are filtered in the remaining single journal portions around 80% of the papers are kept. For comparison, if one excluded all papers from the dataset that had at least one author whose given name is only provided in initialized form not only would over half of the dataset be excluded, also would this filtering be biased against large collaborations since there name abbreviation is used to save printing space.

3. Materials and Methods

The citation data used in this study is composed of all APS journals published between July 1893 and December 2009 and the associated information about citations between these papers and may therefore be assumed to cover many major $20th$ century contributions to physics[1]. The data is provided in separate files corresponding to the individual APS journals which emphasize different topics; we chose to exclude *Reviews of Modern Physics* and the *online-only* journals *Physical Review Special Topics - Accelerators and Beams* and Physics Education Research due to there considerable smaller size compared to the remaining journals. These remaining journals vary in size from *Physical Review E* with 35022 to *Physical Review B* with 133269 publications. In addition to extensive bibliographic metadata like authorship, publication date and actual APS journal the article has been published in the dataset also provides *article type* tags, thereby allowing us to filter non-standard material (i.e. those tagged *comment, erratum, reply, editorial, essay, publisher-note, retraction, miscellaneous*) and restrict our analysis to standard research publications.

Since motif enumeration is computationally costly, for most of the calculations the RANDESU algorithm as described by Wernicke [33] is applied which instead of enumerating all subgraphs samples uniformly from the set of all motifs. By performing duplicate runs of the motif score computation with the same parameters we checked that the scores are not sensitive to the sampling procedure. For smaller networks the full enumeration is performed by application of the ESU algorithm, presented also in [33].

4. Results

We computed the q_m score in both the single journal portions and the whole APS dataset and found that the shuffling of the citation frequencies among the papers indeed yields a uniform distribution of the scores. As mentioned

[1]The dataset can be requested at the American Physical Society, `https://publish.aps.org/datasets`

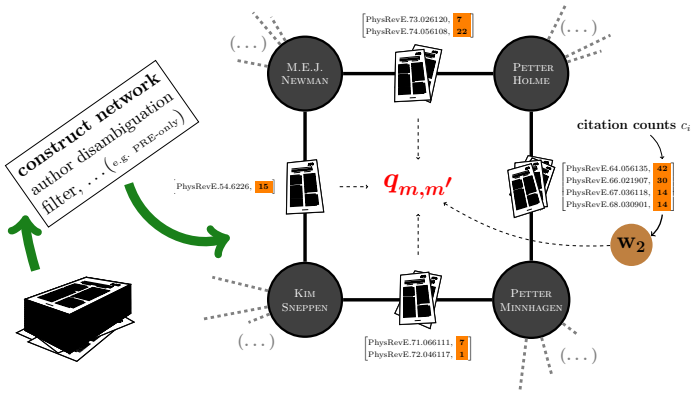

Figure 2: A sketch of the score computation.

above, in all cases the appropriate edge weight to achieve this was $w_2 = \langle c \rangle_e$, i.e. the averaged citation frequency of all papers on the corresponding edge. In the following we show the influence of a variation of the maximal number of authors above which a paper is excluded from the data, as well as the sensitivity of the q_m score against different implementations of the author disambiguation task.

4.1. Maximum Collaboration Size (MCS) scan

In Fig. 3 the average network degree as a function of the MCS value is shown for both the allINIT and the STRICT disambiguation. Increasing the MCS value translates into the introduction of new potential nodes and edges to the network, and since a publication with a authors can produce up to $a(a-1)/2$ new edges $\langle k \rangle$ grows. Unlike PRA, PRB and PRE, the journals PRC, PRD and PRL which feature larger collaborations do not show a saturated average degree in the shown MCS interval. The fact that the STRICT-disambiguation results in more distinct authors than the allINIT implemantation is reflected by the lower average degrees.

Especially with the STRICT disambiguation scheme we can reproduce the main finding of [1], the higher q_m score for the box motif, for small enough values of the maximum collaboration size. Although in few journal portions this result is rather robust, in most cases the box motif signal tends to decrease with increasing MCS; Fig. 4 shows two examples. While in the merge of the journals box motif keeps the highest score in the shown MCS range, in PRL the signal is lost.

4.2. Influence of the author disambiguation

The exact implementation of the author disambiguation turned out to be crucial considering the q_m distribution. From the network perspective it can be interpreted as a local perturbation (as illustrated in the sketch in Fig. 5) which strongly influences the q_m scores which are computed on the few-node subgraph

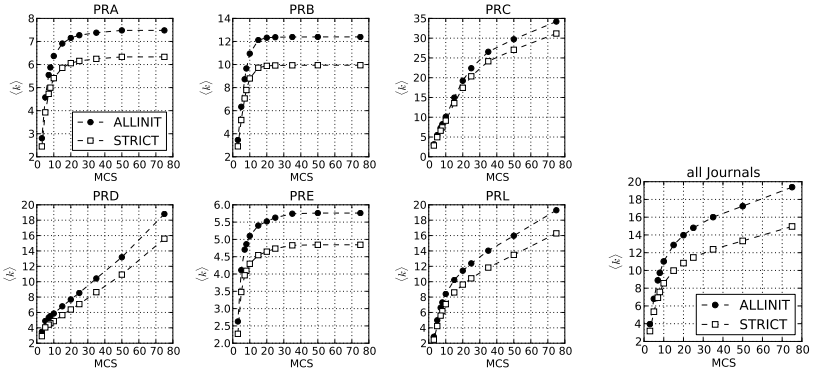

Figure 3: A scan of the MCS value translates to a variation of the average degree.

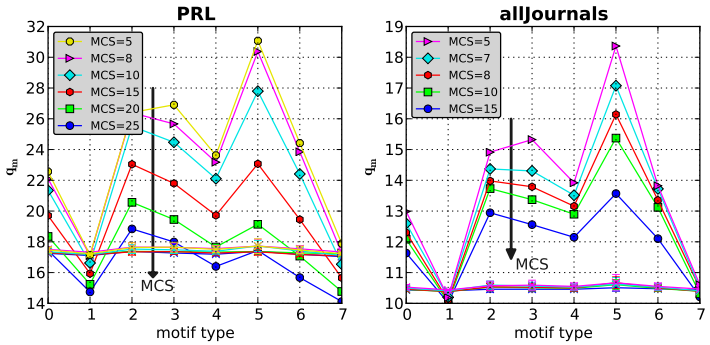

Figure 4: Two example MCS scans (STRICT disambiguation). Weaker signal for larger MCS.

scale. The journal merge does not show the signal any longer and also in the case of PRB and PRE shown in Fig. 5 the box motif scores are suppressed.

The results for the SPLIT disambiguation (which has only been applied for MCS= 10) are shown in Fig. 6. Only PRA and PRL do not show a maximum score for the box motif.

4.3. Distributions of the $q_{m,m'}$ scores

The q_m are averaged values over the $q_{m,m'}$ scores of the single instances of the specific motif type m and as such are influenced of course not only by the top ranked but also by the motif instances with lowest scores. In Fig. 7 the distributions of the $q_{m,m'}$ values for the different motif types, m, are shown for the example of the PRB network portion. In the depicted range the distributions can be grouped according to the minimal number of distinct papers required to construct the specific motif (bottom row in Fig. 1). In the box motif which requires the largest minimum number of distinct papers very low $q_{m,m'}$ scores are suppressed while the three- and four-node cliques which need only one mutual

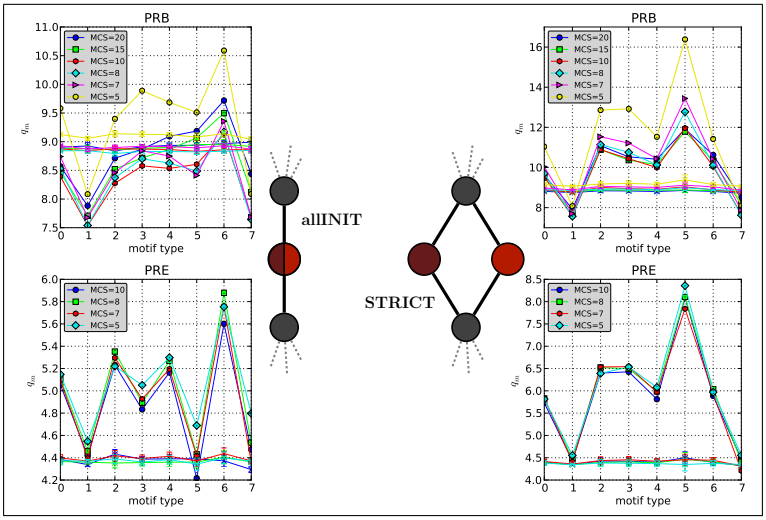

Figure 5: The influence of the local topological perturbation introduced by switching between the allINIT and the STRICT disambiguation scheme for the PRB and the PRE single journal portions.

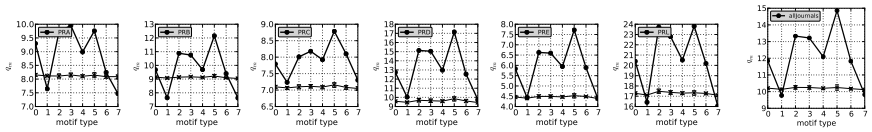

Figure 6: The q_m scores for for the single journal portions and the journal merge with the SPLIT disambiguation applied.

publications among the authors often show very small scores. This is due to the distribution of the citation count used in the computation of the scores: the many poorly cited publications translate directly to low scores for cliques (which do not contain many other publications), while it is unlikely for the at least four edges of a box motif to contain only poorly cited papers. Indeed, the $q_{m,m'}$ distributions of the motif types $0, 4, 6$ and $2, 3$ which require the intermediate number of 2 and 3 edges, respectively, is consistent with this interpretation.

A paper with a authors contributes to $a(a-1)/2$ edges which in turn may contribute to very many motif instances. One corresponding effect can be seen when examining the top motif instances according to the $q_{m,m'}$ scores: These lists can be dominated by highly-cited publications. For example, more than half of the top motifs in the PRA journal portion (SPLIT disambiguation, MCS=10) share the publication PhysRevA.57.120 (D.P. DiVincenzo, D. Loss: *Quantum computation with quantum dots*). If one is not, however, interested in compiling lists of top instances, but only in averaged values like the q_m, another approach might be used. The q_m for a specific motif type m can be rephrased as a weighted sum over the citation counts of all papers that contribute to instances

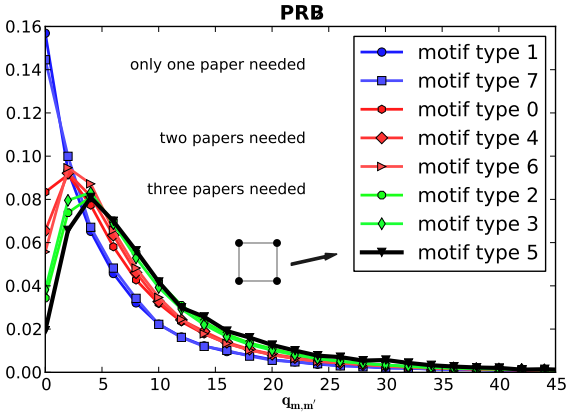

Figure 7: Normalized frequency of the $q_{m,m'}$ scores for the different motif types. The PRB journal portion has been used with MCS= 10 and the SPLIT filtering has been applied.

of that motif; the weights, however, depend not only on how often the paper contributes but also on with how many other papers it shared the edges. A simple alternative would be to set these weights to unity, i.e. compute the average citation frequency of the papers that contribute at least to one of the instances of the motif under consideration. The only information needed from the motif enumeration would then be if a given paper has contributed to an instance of a specific motif type. Fig. 8 shows the example of the PRA data. (We checked via sampling from the citation distribution that high value for the box motif is not simply due to the smaller sample size.) Indeed, also for the other journal portions the box motif seems to be constructed of on average higher cited publications.

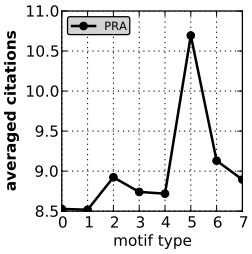

Figure 8: The average citation count of the publications contributing to at least one instance of a given motif type for the PRA data (SPLIT disambiguation).

5. Conclusions

We applied a scheme, originally proposed by Krumov *et al.* [1], which evaluates few-node subgraphs in a collaboration network according to the number of citations the publications in the corresponding subgraph instances received. It turns out that this score is sensitive especially to the local perturbations introduced by changing the implementation of the author disambiguation task and should therefore be used with care. However, the average citation frequency of the papers contributing to a specif subgraph type might strengthen the claim that the collaboration pattern given represented by the box motif is to some extent correlated with higher impact of the resulting papers.

References

[1] L. Krumov, C. Fretter, M. Müller-Hannemann, K. Weihe, M.-T. Hütt, Motifs in co-authorship networks and their relation to the impact of scientific publications, The European Physical Journal B-Condensed Matter and Complex Systems 84 (4) (2011) 535540.

[2] T. S. Kuhn, The Structure of Scientific Revolutions, ISSR library, University of Chicago Press, 1996.

[3] D. J. D. S. Price, Networks of Scientific Papers, Science 149 (1965) 510–515.

[4] R. Albert, A.-L. Barabási, Statistical mechanics of complex networks, Rev. Mod. Phys. 74 (2002) 47–97. doi:10.1103/RevModPhys.74.47.
URL http://link.aps.org/doi/10.1103/RevModPhys.74.47

[5] D. J. D. S. Price, A General Theory of Bibliometric and Other Cumulative Advantage Processes, J. Amer. Soc. Inform. Sci. 27 (1976) 292–306.

[6] A.-L. Barabsi, R. Albert, Emergence of scaling in random networks, Science 286 (5439) (1999) 509–512. doi:10.1126/science.286.5439.509.

[7] S. N. Dorogovtsev, J. F. F. Mendes, A. N. Samukhin, Structure of Growing Networks with Preferential Linking, Phys. Rev. Lett. 85 (2000) 4633–4636. doi:10.1103/PhysRevLett.85.4633.

[8] S. N. Dorogovtsev, J. F. F. Mendes, Scaling properties of scale-free evolving networks: Continuous approach, Phys. Rev. E 63 (2001) 056125. doi:10.1103/PhysRevE.63.056125.

[9] S. N. Dorogovtsev, J. F. F. Mendes, Evolution of networks with aging of sites, Phys. Rev. E 62 (2000) 1842–1845. doi:10.1103/PhysRevE.62.1842.

[10] G. Bianconi, A.-L. Barabási, Bose-Einstein condensation in complex networks, Phys. Rev. Lett. 86 (2001) 5632–5635.

[11] P. L. Krapivsky, S. Redner, Organization of growing random networks, Phys. Rev. E 63 (2001) 066123. doi:10.1103/PhysRevE.63.066123.

[12] P. L. Krapivsky, S. Redner, F. Leyvraz, Connectivity of growing random networks, Phys. Rev. Lett. 85 (2000) 4629–4632. doi:10.1103/PhysRevLett.85.4629.

[13] F. Radicchi, S. Fortunato, C. Castellano, Universality of citation distributions: Toward an objective measure of scientific impact, Proceedings of the National Academy of Sciences 105 (45) (2008) 17268–17272. doi:10.1073/pnas.0806977105.

[14] J. Bollen, H. Van de Sompel, A. Hagberg, R. Chute, A Principal Component Analysis of 39 Scientific Impact Measures, PLoS ONE 4 (6) (2009) e6022. doi:10.1371/journal.pone.0006022.

[15] P. Chen, H. Xie, S. Maslov, S. Redner, Finding scientific gems with Googles PageRank algorithm, Journal of Informetrics 1 (1) (2007) 8 – 15. doi:10.1016/j.joi.2006.06.001.

[16] D. Walker, H. Xie, K.-K. Yan, S. Maslov, Ranking scientific publications using a model of network traffic, Journal of Statistical Mechanics: Theory and Experiment 2007 (06) (2007) P06010.
URL http://stacks.iop.org/1742-5468/2007/i=06/a=P06010

[17] D. F. Klosik, S. Bornholdt, The citation wake of publications detects nobel laureates' papers, arXiv:1301.7471 [physics].
URL http://arxiv.org/abs/1301.7471

[18] J. E. Hirsch, An index to quantify an individual's scientific research output, Proceedings of the National Academy of Sciences 102 (2005) 16569–16572.

[19] F. Radicchi, S. Fortunato, B. Markines, A. Vespignani, Diffusion of scientific credits and the ranking of scientists, Phys. Rev. E 80 (2009) 056103. doi:10.1103/PhysRevE.80.056103.

[20] J. D. West, T. C. Bergstrom, C. T. Bergstrom, The Eigenfactor MetricsTM: A network approach to assessing scholarly journals, College & Research Libraries 71 (3) (2010) 236–244.

[21] P. Chen, S. Redner, Community structure of the Physical Review citation network, Journal of Informetrics 4 (3) (2010) 278 – 290. doi:10.1016/j.joi.2010.01.001.

[22] Z.-X. Wu, P. Holme, Modeling scientific-citation patterns and other triangle-rich acyclic networks, Phys. Rev. E 80 (2009) 037101. doi:10.1103/PhysRevE.80.037101.

[23] M. E. J. Newman, Scientific collaboration networks.i. network construction and fundamental results, Physical Review E 64 (1) (2001) 016131. doi:10.1103/PhysRevE.64.016131.
URL http://link.aps.org/doi/10.1103/PhysRevE.64.016131

[24] M. E. J. Newman, Scientific collaboration networks. II. shortest paths, weighted networks, and centrality, Physical Review E 64 (1) (2001) 016132. doi:10.1103/PhysRevE.64.016132.
URL http://link.aps.org/doi/10.1103/PhysRevE.64.016132

[25] T. Martin, B. Ball, B. Karrer, M. E. J. Newman, Coauthorship and citation patterns in the physical review, Physical Review E 88 (1) (2013) 012814. doi:10.1103/PhysRevE.88.012814.
URL http://link.aps.org/doi/10.1103/PhysRevE.88.012814

[26] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, U. Alon, Network motifs: Simple building blocks of complex networks, Science 298 (5594) (2002) 824–827, PMID: 12399590. doi:10.1126/science.298.5594.824.
URL http://www.sciencemag.org/content/298/5594/824

[27] R. Milo, S. Itzkovitz, N. Kashtan, R. Levitt, S. Shen-Orr, I. Ayzenshtat, M. Sheffer, U. Alon, Superfamilies of evolved and designed networks, Science 303 (5663) (2004) 1538–1542, PMID: 15001784. doi:10.1126/science.1089167.
URL http://www.sciencemag.org/content/303/5663/1538

[28] C. Fretter, M. Mller-Hannemann, M.-T. Htt, Subgraph fluctuations in random graphs, Physical Review E 85 (5) (2012) 056119. doi:10.1103/PhysRevE.85.056119.
URL http://link.aps.org/doi/10.1103/PhysRevE.85.056119

[29] M. E. Beber, C. Fretter, S. Jain, N. Sonnenschein, M. Mller-Hannemann, M.-T. Htt, Artefacts in statistical analyses of network motifs: general framework and application to metabolic networks, Journal of The Royal Society Interface 9 (77) (2012) 3426–3435, PMID: 22896565. doi:10.1098/rsif.2012.0490.
URL http://rsif.royalsocietypublishing.org/content/9/77/3426

[30] C. Schulz, A. Mazloumian, A. M. Petersen, O. Penner, D. Helbing, Exploiting citation networks for large-scale author name disambiguation, arXiv:1401.6157 [physics].
URL http://arxiv.org/abs/1401.6157

[31] T. C. Silva, D. R. Amancio, Network-based stochastic competitive learning approach to disambiguation in collaborative networks, arXiv:1303.0344.
URL http://arxiv.org/abs/1303.0344

[32] S. Milojevi, Accuracy of simple, initials-based methods for author name disambiguation, Journal of Informetrics 7 (4) (2013) 767 – 773. doi:http://dx.doi.org/10.1016/j.joi.2013.06.006.
URL http://www.sciencedirect.com/science/article/pii/S1751157713000539

[33] S. Wernicke, Efficient detection of network motifs, IEEE/ACM Transactions on Computational Biology and Bioinformatics 3 (4) (2006) 347–359. doi:10.1109/TCBB.2006.51.